

Sketching for Simultaneously Sparse and Low-Rank Covariance Matrices

Sohail Bahmani and Justin Romberg
 School of Electrical and Computer Engineering
 Georgia Institute of Technology
 Atlanta, GA 30332
 Email: {sohail.bahmani,jrom}@ece.gatech.edu

Abstract—We introduce a technique for estimating a structured covariance matrix from observations of a random vector which have been sketched. Each observed random vector x_t is reduced to a single number by taking its inner product against one of a number of pre-selected vector a_ℓ . These observations are used to form estimates of linear observations of the covariance matrix Σ , which is assumed to be simultaneously sparse and low-rank. We show that if the sketching vectors a_ℓ have a special structure, then we can use straightforward two-stage algorithm that exploits this structure. We show that the estimate is accurate when the number of sketches is proportional to the maximum of the rank times the number of significant rows/columns of Σ . Moreover, our algorithm takes direct advantage of the low-rank structure of Σ by only manipulating matrices that are far smaller than the original covariance matrix.

I. INTRODUCTION

We introduce and analyze a technique for estimating the covariance matrix Σ of N -dimensional random vector x from samples x_1, x_2, \dots, x_Q . We will show that these quantities can be estimated accurately from *sketches* or *compressed measurements* of x . Different methods for sketching covariance matrices that are either Toeplitz, sparse, or low-rank are studied in [1] and [2]. Sketching of simultaneously structured covariance matrices (e.g., low-rank and sparse matrices or low-rank and Toeplitz matrices) was first considered in [3] and [4]. In particular, it is shown in [3] that simultaneously $K \times K$ -sparse and rank- R covariance matrices can be recovered from $O(K^2 R \log N)$ generic rank-one sketches through minimization of a mixture of the trace norm and the ℓ_1 norm. It is recognized in [3] that the achieved sample complexity is suboptimal. In this paper, we show how the sample complexity can be improved using specifically tailored rank-one sketches. We also demonstrate how the estimation can be performed by manipulating matrices of dimension $N \times R$, where R is the rank of the target.

As the samples x_t are presented, they are mapped into scalar values by taking an inner product against one of L different vectors a_1, \dots, a_L . If a fixed vector a_ℓ is used on T different samples, then we have the estimate

$$a_\ell^\top \Sigma a_\ell = \mathbb{E} [a_\ell^\top x x^\top a_\ell] \approx \frac{1}{T} \sum_{t=1}^T |\langle x_t, a_\ell \rangle|^2. \quad (1)$$

The Q data points are thus turned into L measurements of Σ of the form

$$y[\ell] = a_\ell^\top \Sigma a_\ell + \text{noise}. \quad (2)$$

The measurements $y[\ell]$ can be formed in a decentralized manner, and since they are scalars they are easy to store and communicate. We will show that by choosing the vectors a_ℓ attentively, we can estimate Σ from $L = O(KR \log N)$ rank-one sketches that is much smaller than the number of entries in the covariance matrix, $L \ll N^2$. The proposed sketching scheme is similar to the efficient measurement scheme recently proposed for compressive phase retrieval [5], [6]. Moreover, our algorithm for estimating Σ works by manipulating matrices in factored form, making it scalable for large N regimes.

A. Data model

The data points $x_1, \dots, x_Q \in \mathbb{R}^N$ are independent realizations of a zero-mean random vector with covariance matrix Σ . We will consider the case where Σ is simultaneously sparse and low-rank. This means that we can closely approximate Σ with the factorization

$$\Sigma \approx U U^\top,$$

where U is a $N \times R$ matrix with at most K non-zero rows (we will assume $R \leq K \ll N$). Σ itself can be well-approximated by a matrix with K^2 non-zero terms — all but K rows (or columns) will be approximately zero, and in each of these K significant rows, all but K entries will be approximately zero.

Other than these properties of the covariance matrix, our framework does not depend heavily on the particulars of the distribution of the x_q . The algorithms below depend on having bounds on the approximation error in (1); these bounds might be derived from other known properties of Σ (i.e. its trace or the dynamic range of its eigenvalues) or of the distribution of x .

B. Choosing the sketching vectors

Each data vector x_t is compressed into a single number by taking an inner product against one of L different $a_\ell \in \mathbb{R}^N$. We will denote the set of indices that use vector a_ℓ as \mathcal{T}_ℓ . It is not critically important how the Q observation are divided among the \mathcal{T}_ℓ other than that the sets should have close to the same size. For simplicity, we will assume that all of the index

sets have the same number of terms, $T := |\mathcal{T}_\ell|$ for all ℓ , and so $Q = TL$.

The guarantees for the estimation algorithm presented below depend on the \mathbf{a}_ℓ having certain properties. We restrict the \mathbf{a}_ℓ to lie in an M dimensional subspace of \mathbb{R}^N , generating them using

$$\mathbf{a}_\ell = \Psi^\top \mathbf{w}_\ell, \quad (3)$$

where Ψ is an $M \times N$ matrix whose rows form a basis for this subspace. We will take the \mathbf{w}_ℓ to be randomly generated. The analysis below holds when the \mathbf{w}_ℓ are independent and distributed $\text{Normal}(\mathbf{0}, \mathbf{I})$. However, it is likely that the analysis could be generalized to other distributions, and the only thing we rely on algorithmically is that all L vectors lie in a subspace as in (3).

Our analysis also requires that the matrix Ψ is a stable embedding for $2K$ -sparse vectors. We will assume below that Ψ obeys the *restricted isometry property* [7]

$$(1 - \delta_K) \|\mathbf{z}_1 - \mathbf{z}_2\|_2^2 \leq \|\Psi(\mathbf{z}_1 - \mathbf{z}_2)\|_2^2 \leq (1 + \delta_K) \|\mathbf{z}_1 - \mathbf{z}_2\|_2^2, \quad (4)$$

for all K -sparse $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^N$. There are many examples of matrices that have this property for a number of rows proportional to the sparsity K times a logarithmic factor; see [8], [9] for detailed overviews. If Ψ is populated with independent *subgaussian* random variables, then (4) holds with high probability for $M \gtrsim K \log(N/K)$. More structured matrices that allow for fast computations also have this property. For example, if Ψ consists of M randomly selected rows of an $N \times N$ Fourier matrix, then (4) holds with high probability for $M \gtrsim K \log^4(N)$.

With this subspace conditions, we can write the covariance measurements as

$$y[\ell] = \frac{1}{T} \sum_{t \in \mathcal{T}_\ell} |\langle \mathbf{x}_t, \mathbf{a}_\ell \rangle|^2 \approx \mathbf{w}_\ell^\top \Psi \Sigma \Psi^\top \mathbf{w}_\ell = \langle \Psi \Sigma \Psi^\top, \mathbf{w}_\ell \mathbf{w}_\ell^\top \rangle. \quad (5)$$

That is, each $y[\ell]$ is the matrix inner product of $\Psi \Sigma \Psi^\top$ with an $M \times M$ rank-1 random matrix. This gives the covariance matrix measurements a *nested structure*. We can write

$$\mathbf{y} = \mathcal{W}(\Psi \Sigma \Psi^\top) + \text{noise}.$$

The $N \times N$ covariance matrix Σ is first mapped to an $M \times M$ matrix by applying Ψ to either side. Then \mathcal{W} maps this result to \mathbb{R}^L by taking the series of matrix inner products in (5).

In the next section, we will see how this nested structure allows us to decouple the estimation process into a low-rank estimation stage followed by two sparse approximation stages.

In Section II, we present results that relate the number of sketches L and their accuracy (which is controlled by T) to the accuracy of our estimation procedure.

C. Estimating the covariance

Let $\Sigma^* \in \mathbb{R}^{N \times N}$ be a rank- R positive semidefinite matrix that has at most K nonzero rows (and columns) with $R < L \ll N$. Given a matrix $\Psi \in \mathbb{R}^{M \times N}$ and a linear operator \mathcal{W} :

$\mathbf{B} \mapsto [\mathbf{w}_\ell^\top \mathbf{B} \mathbf{w}_\ell]_{\ell=1}^L$, where $\mathbf{w}_\ell \sim \text{Normal}(\mathbf{0}, \mathbf{I})$, we consider the problem of estimating Σ^* from noisy linear measurements of the form

$$\mathbf{y} = \mathcal{W}(\Psi \Sigma^* \Psi^\top) + \mathbf{z}, \quad (6)$$

where $\mathbf{z} \in \mathbb{R}^L$ is the noise vector that is bounded as $\|\mathbf{z}\|_2 \leq \varepsilon$.

We propose the following two-stage procedure for estimation of Σ^* .

- 1) **Low-rank estimation stage:** Since Σ^* is low-rank, we know that $\Psi \Sigma \Psi^\top$ will be low-rank as well. Our first stage “inverts” the $\mathcal{W}(\cdot)$ operator by looking for a low-rank matrix that explains the measurements \mathbf{y} . There are a number of ways we might do this, but here we will use the standard convex relaxation that minimizes the trace norm (nuclear norm) in place of the rank:

$$\begin{aligned} \hat{\mathbf{B}} \in \underset{\mathbf{B} \succeq \mathbf{0}}{\text{argmin}} \quad & \text{trace}(\mathbf{B}) \\ \text{subject to} \quad & \|\mathcal{W}(\mathbf{B}) - \mathbf{y}\|_2 \leq \varepsilon. \end{aligned} \quad (7)$$

The output $\hat{\mathbf{B}}$ of this first stage can be thought of as an estimate of $\Psi \Sigma^* \Psi^\top$.

- 2) **Sparse estimation stage:** In the second stage, we invert the action of Ψ on the left and Ψ^\top on the right. It is conceivable that these two steps could be combined into a single sparse approximation step, but performing them sequentially leads to a natural analysis (as we will see in Section II).

- a) Since we are given $\hat{\mathbf{B}} \approx \Psi \Sigma^* \Psi^\top$, we start by looking for a matrix with a small number of non-zero rows whose range is close to the range of Σ^* . The fact that $\hat{\mathbf{B}}$ is approximately rank R allows us to work with a rank R approximation of $\hat{\mathbf{B}}$ and be computationally efficient. We start by computing the $M \times R$ matrix $\hat{\mathbf{U}}$ of the top R (unit-norm) eigenvectors of $\hat{\mathbf{B}}$. We can then look for a row-wise sparse $N \times R$ matrix that is close to $\hat{\mathbf{B}} \hat{\mathbf{U}}$ after an application of Ψ on the left:

$$\begin{aligned} \hat{\mathbf{Q}}_1 \in \underset{\mathbf{Q}}{\text{argmin}} \quad & \|\mathbf{Q}\|_{1,2} \\ \text{subject to} \quad & \left\| \Psi \mathbf{Q} - \hat{\mathbf{B}} \hat{\mathbf{U}} \right\|_F \leq \frac{c_1 \varepsilon}{\sqrt{L}}. \end{aligned} \quad (8)$$

The $\|\cdot\|_{1,2}$ norm above is a convex relaxation for the number of non-zero rows; $\|\mathbf{Z}\|_{1,2} = \sum_{n=1}^N \|\mathbf{Z}_{n,:}\|_2$, where $\mathbf{Z}_{n,:}$ is the n th row of \mathbf{Z} . Conceptually (and this is formalized in the analysis below), the output $\hat{\mathbf{Q}}_1$ can be used to form the approximation $\hat{\mathbf{Q}}_1 \hat{\mathbf{U}}^\top \approx \Sigma^* \Psi^\top$. Thus we have effectively undone the action of Ψ on the left of Σ^* .

- b) The second sparse approximation step does a similar computation on the row-space of the output above. We compute $\hat{\mathbf{V}} \in \Sigma^{N \times R}$, the matrix of

the left singular vectors of \hat{Q}_1 , and compute

$$\begin{aligned} \hat{Q}_2 \in \operatorname{argmin}_Q \|\mathbf{Q}\|_{1,2} \\ \text{subject to } \left\| \Psi \mathbf{Q} - \hat{\mathbf{U}} \hat{Q}_1^\top \hat{\mathbf{V}} \right\|_F \leq \frac{c_2 \varepsilon}{\sqrt{L}}. \end{aligned} \quad (9)$$

Conceptually, the output \hat{Q}_2 undoes the action of Ψ^\top on the right, giving us the approximation $\hat{Q}_2 \hat{\mathbf{V}}^\top \approx \Sigma^*$

The output of the second stage are two $N \times R$ matrices $\hat{\mathbf{V}}$ and \hat{Q}_2 . These can be used to form an approximate factorization of Σ^* . In Section II below, we give a mathematical guarantee on how close

$$\hat{\Sigma} = \hat{Q}_2 \hat{\mathbf{V}}^\top, \quad (10)$$

is to the true covariance Σ^* .

The constants c_1 and c_2 in the sparse estimation stage above are user-defined; the analysis below relies on a particular choice for these parameters.

The algorithm above is carefully designed to be sublinear in the number of entries in the $N \times N$ matrix Σ^* . Stage 1 above is an optimization program over the cone of $M \times M$ SDP matrices, and we have seen that we can take $M \gtrsim K \log^\alpha N$. The sparse approximation stage exclusively handles $N \times R$ matrices. This allows our algorithmic framework to scale to regimes where the dimension N and the number of samples Q are large.

D. Noise magnitude

In general, all the algorithm proposed above needs is a bound on the total size of the error in the measurements of the covariance matrix. With the model in (1) and (2), this error is simply the deviation of a quadratic function of the random vector \mathbf{x} from its mean. If we have information about the distribution of \mathbf{x} , we may be able to derive the desired bound in a principled way.

To demonstrate this, suppose that $\mathbf{x} \sim \text{Normal}(\mathbf{0}, \Sigma)$ which implies $\langle \mathbf{x}, \mathbf{a}_\ell \rangle \sim \text{Normal}(0, \mathbf{a}_\ell^\top \Sigma \mathbf{a}_\ell)$. This means that

$$y[\ell] = \frac{1}{T} \sum_{t \in \mathcal{T}_\ell} |\langle \mathbf{x}_t, \mathbf{a}_\ell \rangle|^2$$

is proportional to a Chi-squared random variable with T degrees of freedom. Since $\mathbb{E}[y[\ell]] = \mathbf{a}_\ell^\top \Sigma \mathbf{a}_\ell$, the mean energy of the ℓ th component $e[\ell] := y[\ell] - \mathbf{a}_\ell^\top \Sigma \mathbf{a}_\ell$ of the error is

$$\mathbb{E}[e[\ell]^2] = \frac{2(\mathbf{a}_\ell^\top \Sigma \mathbf{a}_\ell)^2}{T},$$

and the total error $\|e\|_2^2$ can be shown to concentrate around

$$\mathbb{E}[\|e\|_2^2] = \frac{2}{T} \sum_{\ell=1}^L (\mathbf{a}_\ell^\top \Sigma \mathbf{a}_\ell)^2.$$

The obtained concentration bounds may depend on the (a priori unknown) covariance Σ , but usually they can be approximated by some attributes of Σ . If the \mathbf{a}_ℓ are chosen as in Section I-B, $\mathbf{a}_\ell^\top \Sigma \mathbf{a}_\ell$ will itself be a weighted sum of Chi-squared random variables whose moments can be calculated in terms of the Frobenius and spectral norms of Σ .

II. ANALYSIS

Our main theorem shows that for an appropriate number of sketches, the estimation algorithm detailed in Section I-C produces a provably good estimate of Σ^* . In addition to the number of sketches L being sufficiently large, we also assume that Ψ obeys the restricted isometry property in (4).

Theorem 1. *There is a constant C_1 such that if*

$$L \geq C_1 \cdot R \cdot M, \quad (11)$$

then for appropriately chosen constants c_1 and c_2 in (8) and (9) above, with probability exceeding $1 - e^{-O(M)}$ the estimate in (10) obeys

$$\left\| \hat{\Sigma} - \Sigma^* \right\|_F \leq \frac{C_2 \varepsilon}{\sqrt{L}},$$

where C_2 is another constant which depends on c_1, c_2 , and the restricted isometry constant δ_K .

The theorem above gives us a uniform guarantee; it holds simultaneously for all rank- R and row-wise K -sparse Σ^* for the same $\{\mathbf{a}_\ell\}$. We will sketch a proof below, withholding some of the details due to space constraints.

The accuracy of the first stage is established through a simple application of a recent result in the theory of low-rank matrix recovery. The work [10] establishes uniform bounds on the recovery of low-rank matrices from inner products against a series of independent, rank-1 symmetric random matrices. This exactly describes the $\mathcal{W}(\cdot)$ operator in (6). A direct application of the main theorem in that paper shows that for L as in (11), we have

$$\left\| \hat{\mathbf{B}} - \Psi \Sigma^* \Psi^\top \right\|_F \leq \frac{c\varepsilon}{\sqrt{L}}, \quad (12)$$

with high probability for an absolute constant $c > 0$.

The accuracy of the second stage follows from the lemma below, which we will prove at the end of the section.

Lemma 1. *Suppose that for some rank- R and row-wise K -sparse matrix Σ^\sharp , a matrix \mathbf{B}^\sharp , and a constant ϵ we have*

$$\left\| \Psi \Sigma^\sharp - \mathbf{B}^\sharp \right\|_F \leq \epsilon,$$

where Ψ obeys (4) with a sufficiently small δ_K . Let \mathbf{U}^\sharp denote the top R right singular vectors of \mathbf{B}^\sharp and

$$\mathbf{Q}^\sharp = \operatorname{argmin}_Q \|\mathbf{Q}\|_{1,2} \quad (13)$$

$$\text{subject to } \left\| \Psi \mathbf{Q} - \mathbf{B}^\sharp \mathbf{U}^\sharp \right\|_F \leq 2\epsilon.$$

Then we have

$$\left\| \mathbf{Q}^\sharp \mathbf{U}^{\sharp\top} - \Sigma^\sharp \right\|_F \leq C\epsilon,$$

for some constant $C > 0$ depending only on δ_K .

With $\hat{\mathbf{U}}$ as the top R eigenvectors of the output of the first stage $\hat{\mathbf{B}}$, Lemma 1 tells us that the result \hat{Q}_1 of solving (8) in Stage 2a will obey

$$\left\| \hat{\mathbf{U}} \hat{Q}_1^\top - \Psi \Sigma^* \right\|_F = \left\| \hat{Q}_1 \hat{\mathbf{U}}^\top - \Sigma^* \Psi^\top \right\|_F \leq \frac{c'\varepsilon}{\sqrt{L}},$$

for some constant c' . For stage 2b, with $\hat{\mathbf{V}} \in \mathbb{R}^{N \times R}$ as the top R left singular vectors of $\hat{\mathbf{Q}}_1$, we can again invoke the lemma and conclude that for some constant C_2 we have

$$\|\hat{\Sigma} - \Sigma^*\|_F = \|\hat{\mathbf{Q}}_2 \hat{\mathbf{V}}^\top - \Sigma^*\|_F \leq \frac{C_2 \varepsilon}{\sqrt{L}}.$$

It remains to prove Lemma 1. Because $\Psi \Sigma^\sharp$ is rank- R and $B^\sharp U^\sharp U^{\sharp\top}$ is the best rank- R approximation of B^\sharp , we have

$$\begin{aligned} \|\Psi \Sigma^\sharp - B^\sharp U^\sharp U^{\sharp\top}\|_F &\leq \|\Psi \Sigma^\sharp - B^\sharp\|_F + \|B^\sharp - B^\sharp U^\sharp U^{\sharp\top}\|_F \\ &\leq 2 \|\Psi \Sigma^\sharp - B^\sharp\|_F \leq 2\epsilon. \end{aligned}$$

Therefore, we can write

$$\begin{aligned} &\|\Psi \Sigma^\sharp (\mathbf{I} - U^\sharp U^{\sharp\top})\|_F^2 + \|(\Psi \Sigma^\sharp - B^\sharp) U^\sharp\|_F^2 \\ &= \|\Psi \Sigma^\sharp (\mathbf{I} - U^\sharp U^{\sharp\top})\|_F^2 + \|(\Psi \Sigma^\sharp - B^\sharp) U^\sharp U^{\sharp\top}\|_F^2 \\ &= \|\Psi \Sigma^\sharp - B^\sharp U^\sharp U^{\sharp\top}\|_F^2 \leq 4\epsilon^2. \end{aligned}$$

In particular, since Σ^\sharp is row-wise K -sparse, from (4) we have

$$\sqrt{1 - \delta_K} \|\Sigma^\sharp (\mathbf{I} - U^\sharp U^{\sharp\top})\|_F \leq \|\Psi \Sigma^\sharp (\mathbf{I} - U^\sharp U^{\sharp\top})\|_F \leq 2\epsilon. \quad (14)$$

Moreover, we have already obtained

$$\|\Psi \Sigma^\sharp U^\sharp - B^\sharp U^\sharp\|_F \leq 2\epsilon.$$

It follows from this latter bound that $\Sigma^\sharp U^\sharp$ is feasible in (13).

Since Ψ obeys (4), we can use standard results for compressive sensing of block-sparse signals (e.g. [11]) to guarantee that for some absolute constant $c > 0$ we have

$$\|Q^\sharp - \Sigma^\sharp U^\sharp\|_F \leq c\epsilon.$$

Therefore, using (14) and with $C = \sqrt{c^2 + \frac{4}{1-\delta_K}}$ we have

$$\begin{aligned} \|Q^\sharp U^{\sharp\top} - \Sigma^\sharp\|_F^2 &= \|Q^\sharp U^{\sharp\top} - \Sigma^\sharp U^\sharp U^{\sharp\top}\|_F^2 + \|\Sigma^\sharp (\mathbf{I} - U^\sharp U^{\sharp\top})\|_F^2 \\ &= \|Q^\sharp - \Sigma^\sharp U^\sharp\|_F^2 + \|\Sigma^\sharp (\mathbf{I} - U^\sharp U^{\sharp\top})\|_F^2 \\ &\leq C^2 \epsilon^2, \end{aligned}$$

and thereby $\|Q^\sharp U^{\sharp\top} - \Sigma^\sharp\|_F \leq C\epsilon$.

III. SIMULATIONS

We ran numerical simulations on synthetic data as follows. With $N = 1000$ and for $R \in \{2, 4, 8\}$ and $K \in \{10, 11, \dots, 19\}$, in each of the 100 trials we generated an $N \times R$ matrix \mathbf{U} by drawing a $K \times R$ random matrix with iid standard Gaussian entries, modulating its columns by iid Uniform[0, 1], and interleaving its rows with $N - K$ all-zero rows uniformly at random. Then $\Sigma^* = \mathbf{U} \mathbf{U}^\top$ is selected as the target covariance matrix. We also set $M = \lceil 2K (1 + \log \frac{N}{K}) \rceil$ and $L = 3RM$. To compute the measurements we draw 2500 iid samples \mathbf{x}_t from Normal(0, Σ^*) and L samples of $\mathbf{a}_\ell = \Psi^\top \mathbf{w}_\ell$ where the $M \times N$ matrix Ψ is populated with




Fig. 1. The empirical 0.9 quantile of estimation error vs. sparsity K for rank R in $\{2, 4, 8\}$

iid Normal(0, $\frac{1}{M}$) and $\mathbf{w}_\ell \sim \text{Normal}(\mathbf{0}, \mathbf{I})$. In each trial, we applied the proposed method with $\varepsilon = 2RK/\sqrt{T}$, $c_1 = \sqrt{3}$, and $c_2 = 3$. Figure 1 illustrates the 0.9 quantile of the relative error of the estimated covariance matrix as K varies between 10 and 19 for $R = 2, 4$ and 8. The error is almost flat as a function of K , in agreement with the theoretical analysis for the prescribed number of measurements. Higher variations for smaller values of R is the effect of smaller sample size.

REFERENCES

- [1] J. Bioucas-Dias, D. Cohen, and Y. Eldar, "Covalsa: Covariance estimation from compressive measurements using alternating minimization," in *Signal Processing Conference (EUSIPCO), 2014 Proceedings of the 22nd European*, Sept 2014, pp. 999–1003.
- [2] G. Dasarthy, P. Shah, B. Bhaskar, and R. Nowak, "Sketching sparse matrices, covariances, and graphs via tensor products," *Information Theory, IEEE Transactions on*, vol. 61, no. 3, pp. 1373–1388, March 2015.
- [3] Y. Chen, Y. Chi, and A. Goldsmith, "Estimation of simultaneously structured covariance matrices from quadratic measurements," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, May 2014, pp. 7669–7673.
- [4] —, "Exact and stable covariance estimation from quadratic sampling via convex programming," *IEEE Trans. Inform. Theory*, vol. 61, no. 7, pp. 4034 – 4059, July 2015.
- [5] M. Iwen, A. Viswanathan, and Y. Wang, "Robust sparse phase retrieval made easy," *Applied and Computational Harmonic Analysis*, 2015, in press. Preprint, arXiv:1410.5295 [math.NA].
- [6] S. Bahmani and J. Romberg, "Efficient compressive phase retrieval with constrained sensing vectors," in *the Twenty-ninth Annual Conference on Neural Information Processing Systems (NIPS)*, Montreal, Canada, Dec. 2015, preprint arXiv:1507.08254 [cs.IT].
- [7] E. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5245, 2006.
- [8] R. Vershynin, "Introduction to the non-asymptotic theory of random matrices," in *Compressed Sensing, Theory and Applications*, Y. Eldar and G. Kutyniok, Eds. Cambridge University Press, 2012, pp. 210–268.
- [9] H. Rauhut, "Compressive sensing and structured random matrices," *Radon Series Comp. Appl. Math.*, vol. 9, pp. 1–92, 2010.
- [10] R. Kueng, H. Rauhut, and U. Terstiege, "Low rank matrix recovery from rank one measurements," *Applied and Computational Harmonic Analysis*, 2015, in press. Preprint arXiv:1410.6913 [cs.IT].
- [11] Y. C. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Inform. Theory*, vol. 55, no. 11, pp. 5302–5316, November 2009.